Punny Captions: Witty Wordplay in Image Descriptions

Arjun Chandrasekaran[1,*] **Devi Parikh**[1] **Mohit Bansal**[2]
[1]Georgia Institute of Technology [2]UNC Chapel Hill
{carjun, parikh}@gatech.edu mbansal@cs.unc.edu

arXiv:1704.08224v1 [cs.CL] 26 Apr 2017

Abstract

Wit is a quintessential form of rich inter-human interaction, and is often grounded in a specific situation (e.g., a comment in response to an event). In this work, we attempt to build computational models that can produce witty descriptions for a given image. Inspired by a cognitive account of humor appreciation, we employ linguistic wordplay, specifically puns. We compare our approach against meaningful baseline approaches via human studies. In a Turing test style evaluation, people find our model's description for an image to be wittier than a human's witty description 55% of the time!

1 Introduction

> "*Wit is the sudden marriage of ideas which before their union were not perceived to have any relation.*"
>
> – Mark Twain

Wit is integral to inter-human interactions. Witty remarks are often contextual, i.e., grounded in a specific situation. Developing computational models that can understand and emulate subtleties of human expression such as contextual humor, is an important step towards making human-AI interaction more natural and more engaging (Yu et al., 2016). For instance, witty chatbots could help relieve stress and increase user engagement by being more personable, human-like, and trustworthy. Bots could automatically post witty comments (or suggest witty remarks) in response to a friend's post on social media, chat, or messaging. In this work, we attempt to tackle the challenging task of producing a witty (pun-based) remark about a given (possibly boring) image. Our approach is inspired by Suls (1972)'s two-stage cognitive account of humor appreciation. According to Suls, a perceiver experiences humor when a stimulus[1] causes an *incongruity*, followed by *resolution*.

(a) **Generated**: a poll (pole) on a city street at night. **Retrieved**: the light knight (night) chuckled. **Human**: the knight (night) in shining armor drove away.

(b) **Generated**: a bare (bear) black bear walking through a forest. **Retrieved**: another reporter is standing in a bare (bear) brown field. **Human**: the bear killed the lion with its bare (bear) hands.

Figure 1: Sample images and witty descriptions from our generation model, retrieval model and a human. The word in the parenthesis is the pun associated with the image. It is provided as a reference to the source of the unexpected pun used in the caption (e.g., bare and poll).

We attempt to introduce an incongruity by using an unexpected word (pun) in the description of the image. Consider the words (poll, knight) used in descriptions in Fig. 1a. The expectations of a perceiver regarding the image (night, traffic light, street, etc.) are disconfirmed by the description, which mentions 'a poll on the city street'. The incongruity is resolved when the perceiver reinterprets the image description with the original pun word associated with the image (e.g., pole, night, in Fig. 1a). This incongruity followed by resolu-

[*]Part of this work was done when AC was an intern with MB at TTI-Chicago and a student at Virginia Tech.

[1]The stimulus can be a joke or a captioned cartoon.

tion may be perceived as witty[2].

We build two computational models based on our approach to produce witty descriptions for an image. The first model generates witty descriptions for an image by modifying an image captioning model to include the specified pun word in the description during inference. The second model retrieves sentences that are relevant to the image, and also contain a pun, from a large corpus of stories (Zhu et al., 2015). We evaluate the wittiness of these image descriptions via human studies. We compare the top-ranked captions from our models against 3 meaningful, qualitatively different baselines. In a Turing-test style evaluation, we compare head-to-head our top-ranked generated description for an image against a human-written witty description that also uses puns.

Our paper makes the following contributions: To the best of our knowledge, this is the first work that tackles the challenging problem of producing a witty natural language remark in an everyday (boring) context. We present two novel models to produce witty (pun-based) captions for a novel (likely boring) image. Our models rely on linguistic wordplay. They use an unexpected pun in an image description during inference/retrieval. Thus, they do not require to be trained with witty captions. Humans vote the descriptions from the top-ranked *generated* captions 'wittier' than a description decoded via regular inference, a human-witty caption that is mismatched for the given image, and a 'punny' description that is ambiguous, and plausibly witty (see Sec 4). In a Turing test-style evaluation, our model's best description for an image is found to be wittier than witty human-written caption 55% of the time[3].

2 Related Work

Theories of verbal humor. General theory of verbal humor (Attardo and Raskin, 1991) characterizes linguistic stimuli that induce humor. As Binsted (1996) notes, however, implementing computational models of this theory requires severely restricting its assumptions.

Puns. A few works have studied the mechanisms due to which puns induce humor. Pepicello and Green (1984) categorize riddles based on the type of linguistic ambiguity that they exploit – phonological, morphological or syntactic. Zwicky and Zwicky (1986) detail the difference between perfect and imperfect puns, i.e., words that are pronounced exactly the same (homophonic) and those pronounced differently (heterophonic). Miller and Gurevych (2015) and Miller (2016) develop methods to better detect and identify puns.

Generating textual humor. JAPE (Binsted and Ritchie, 1997) is a pun-based riddle generating program. Similar to our work, they leverage phonological ambiguity. While our task involves producing free-form responses to a novel stimulus, JAPE only produces stand-alone "canned" jokes. Also similar to our work is HAHAcronym (Stock and Strapparava, 2005), which generates a funny expansion for a given stimulus (acronym). Unlike our work, HAHAcronym is constrained to a single modality (text), and is limited to producing sets of words. In comparison, our approach is applicable to situational humor in real world human interactions since we generate free form natural-language sentences in response to visual stimuli. Petrovic and Matthews (2013) develop an unsupervised model that produces "*I like my X like I like my Y, Z*" jokes.

Generating multi-modal humor. Wang and Wen (2015) predict a meme's text based on a given funny image. Similarly Shahaf et al. (2015) and Radev et al. (2015) learn to rank cartoon captions based on their funniness. Unlike the typically boring context (images) in our task, memes and cartoons involve a context (images) that are already funny or atypical[4]. Chandrasekaran et al. (2016) modify an abstract scene to make it more funny. Their task is restricted to altering the input (visual) modality, while our task is to generate witty natural language remarks for a novel image.

3 Approach

The lack of large corpora of witty remarks and the relationship of wit with surprise, add to the challenge of learning to be witty via purely data driven methods. Our approach employs linguistic wordplay to overcome some of these challenges. We now describe our pipeline and models in detail.

Tags. The first step towards producing a contextually witty remark is to identify concepts that are

[2]The perceiver fails to appreciate wit if the process of 'solving' (resolution) is trivial (the joke is obvious) or too complex (does not 'get' the joke).

[3]Please note that in a Turing test, a machine-approach would equal a human at 50%.

[4]E.g., "LOL-cats" (involving funny cat photos), "Bieber-memes" (involving modified pictures of a Justin Bieber) and abstract cartoons that involve talking animals.

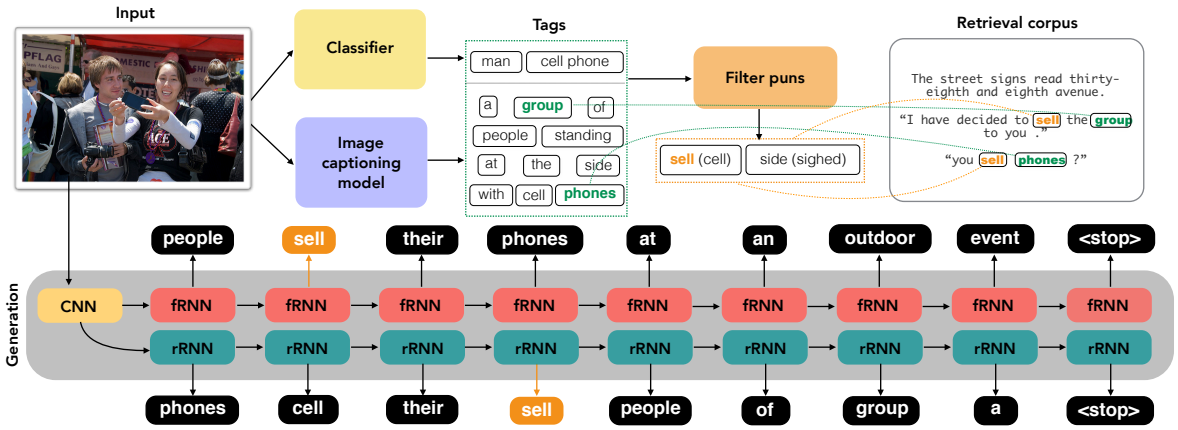

Figure 2: Generating and retrieving descriptions about an image that contain a pun.

relevant to the context (input image). In some cases, an image is directly associated with relevant concepts, e.g., tags posted on social media. We consider the general case where such tags are unavailable, and automatically extract tags associated with an image. First, we recognize objects in the image using an image classifier. We utilize the top K[5] predictions from a state-of-the-art Inception-ResNet-v2 model trained for image classification on ImageNet (Deng et al., 2009).

Second, we consider the words from a (boring) description of an image, generated by the Show-and-Tell (Vinyals et al., 2016) image captioning model. The architecture uses an Inception-v3 CNN encoder (Szegedy et al., 2016) and LSTM decoder, and is trained on the COCO dataset (Lin et al., 2014). As shown in Fig. 2, we combine object labels from the classifier with words from the caption (ignoring stopwords) to produce a set of tags associated with an image.

Puns. We construct a list of heterographic homophones by mining the web. To increase coverage, we use a model from automatic speech recognition research (Jyothi and Livescu, 2014) that predicts the edit distance between two words based on articulatory features. We consider only pairs of words with an edit distance of 0 and manually eliminate false positives. Our list of puns has a total of 1067 unique words (931 from the web and 136 from the speech recognition model).

Pun vocabulary. We utilize our pun list to *filter puns* for a given image, i.e., we identify words associated with an image that have phonologically identical counterparts. The result is a *pun vocabulary* for each image (see Fig. 2).

Generating punny image captions. Based on an image and its *pun vocabulary*[6], we generate witty descriptions (shown in gray in Fig. 2). We use the Show-and-Tell architecture (described above), which decodes the words of a caption conditioned on an image. At specific time-steps during inference (shown in orange in Fig. 2), we force the model to produce a phonological counterpart of a pun word that is associated with the image (in our example, 'sell' or 'sighed'). We achieve this by limiting the vocabulary of the decoder to contain only the counterparts of the image-puns at that time-step. At time-steps that follow, the decoder produces a new word, conditioned on all previously decoded words. Thus, the decoder attempts to produce sentences that flow well based on previously uttered words. A downside to this is that introducing puns at later time-steps results in less grammatical sentences. We overcome this by training two models that decode an image description in forward (start to end) and reverse (end to start) directions, depicted as 'fRNN' and 'rRNN' in Fig. 2 respectively. The forward RNN and reverse RNN generate sentences in which the pun appears in each of the first T and last T positions, respectively[7]. This results in a pool of candidate witty captions. In Fig. 2, a pun is chosen to be decoded in the 2nd and 4th time-steps during inference of the forward and reverse RNN respectively.

Retrieving punny image captions. We attempt to leverage the usage of natural, human-written sentences which are relevant (yet unexpected) in the

[5]In our experiments, we use $K = 5$.

[6]On average, we find that 2 in 5 images have puns associated with them, i.e., have non-empty Pun vocabularies.

[7]In our experiments we choose T=5 and a beam size of 6.

given context. Concretely, we attempt to retrieve natural language sentences from a combination of the Book Corpus (Zhu et al., 2015) and corpora from the NLTK toolkit (Loper and Bird, 2002). We have two constraints on the retrieved sentence. First, to introduce an incongruity, the retrieved sentence must contain the counterpart of the pun word that is associated with the image. Second, to ensure contextual relevance, the retrieved sentence must have *support* in the image, i.e., it must contain at least one image tag. This results in a pool of candidate captions that are perfectly grammatical, a little unexpected, yet relevant to the given context (image).

Ranking. We rank captions in the pools of candidates from *both* models (generation and retrieval), according to log-probability score from the image captioning model. This results in top-ranked descriptions that are both relevant to the image and grammatically correct. We then perform non-maximal suppression, i.e., eliminate captions that are similar[8] to a higher-ranked caption to reduce the pool to a smaller, more diverse set. The top 3 ranked captions from each model are the 'best' captions. The generation model produces better descriptions among our two models.

Data. We produce witty descriptions for images from the validation set of COCO that have puns associated with them. We evaluate them via human studies on a random subset of 100 images.

4 Experiments

Baselines. We compare the wittiness of descriptions generated by our model against 3 qualitatively different baselines. **Regular inference** generates a fluent caption relevant to the image, but is not attempting to be witty. **Witty mismatch** is a human-written witty caption, but for a different image from the one being evaluated. This baseline results in a witty caption, but does not attempt to be relevant to the image. **Ambiguous** is a 'punny' caption where a pun word in the boring (regular) caption is replaced by its counterpart. This caption is likely to contain content that is relevant to the image but it also contains a pun that is not relevant to the image *or* to the rest of the caption.

Annotations. We asked people on AMT to vote for the wittier among the a given pair of descriptions for an image. We compared head-to-head, each of the 4 output captions from our models (3 high scoring, 1 low scoring, described in Sec. 3), against each of the 3 baseline captions. We choose the majority among 9 votes for each relative choice.

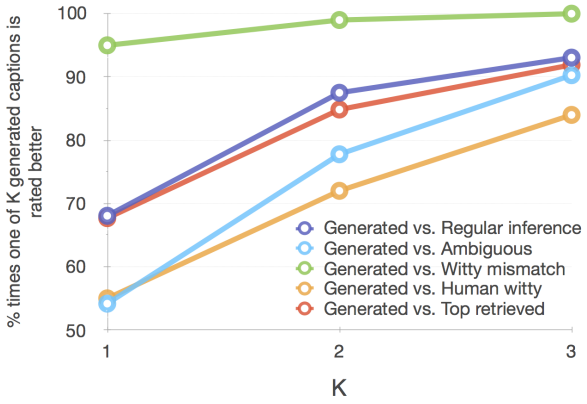

Figure 3: Comparison of wittiness of the top 3 generated captions vs. other approaches. As we increase the number of generated captions (K), recall steadily increases.

Metric. In Fig. 3, we report performance of our model using the Recall@K metric. For each $K = 1, 2, 3$, we compute the percentage of images for which at least one of the K 'best' descriptions from our model outperformed a baseline. As described earlier, our 'best' captions are the top 3 captions, sorted by their log. probability according to our image captioning model.

Quantitative results. As we see in Fig. 3, the image descriptions from our generation approach are voted wittier than all baselines ($>50\%$) even at $K = 1$. As K increases, the recall steadily increases. This indicates that generated captions are witty in the context of the image, and are wittier than a naive approach that introduces ambiguity. The retrieved captions on the other hand are neither witty nor relevant for the image – they are less witty than *Regular inference* and *Ambiguous* descriptions. Further, in a head-to-head comparison, we find the generated captions to be wittier than the retrieved captions (see Fig. 3).

We also validate our choice of ranking captions based on the image captioning model score. We observe that a 'bad' caption, i.e., one that is ranked lower by our model, is significantly less witty than the top 3 output captions.

Human-written witty captions. We ask people on AMT to write a witty description for an image

[8] Two sentences are similar if the cosine similarity between the average of the Word2Vec (Mikolov et al., 2013) representations of words in each sentence is ≥ 0.8.

(a) **Generated**: a bear that is bare (bear) in the water. **Retrieved**: water glistened off her bare (bear) breast. **Human**: you won't hear a creak (creek) when the bear is feasting.

(b) **Generated**: a bored (board) bench sits in front of a window. **Retrieved**: Wedge sits on the bench opposite Berry, bored (board). **Human**: could you please make your pleas (please)!

(c) **Generated**: a woman sell (cell) her cell phone in a city. **Retrieved**: Wright (right) slammed down the phone. **Human**: a woman sighed (side) as she regretted the sell.

(d) **Generated**: a loop (loupe) of flowers in a glass vase. **Retrieved**: the flour (flower) inside teemed with worms. **Human**: piece required for peace (piece).

(e) **Generated**: a female tennis player caught (court) in mid swing. **Retrieved**: my shirt caught (court) fire . **Human**: the woman's hand caught (court) in the center.

(f) **Generated**: broccoli and meet (meat) on a plate with a fork. **Retrieved**: "you mean white folk (fork)". **Human**: the folk (fork) enjoyed the food with a fork.

Figure 4: Sample images and witty descriptions from our generation model, retrieval model and a human. The word in the parenthesis is the pun associated with the image. It is provided as a reference to the source of the unexpected pun used in the caption (e.g., bare/creak, bored, sell/Wright/sighed, loop/flour/peace, caught and meet/folk in captions (a) to (f) respectively).

using a (given) pun associated with it. We then ask a different set of people to compare this description against the top ranked description from our model, in a Turing-test style evaluation. Descriptions from our model are wittier than humans 55% of the time.

Qualitative analysis. We observe that the generation model uses interesting 'techniques' to pro-

duce witty captions. For instance, in Fig. 4a, it employs alliteration using the original pun (bear) and its counterpart (bare). In another example, the caption in Fig 4b makes sense for both the original pun associated with the image (board), and its phonological counterpart (bored). In a few cases, such as in Fig. 4f, the model naively replaces a pun associated with the image (meat) with its counterpart (meet) in a description.

The retrieved sentence often contains words or phrases that are irrelevant to the context of the image, as we see in Fig. 4b. This is a likely reason for why a retrieved sentence containing a pun is perceived as less witty when compared with witty descriptions generated for the image.

5 Discussion

Since wit involves unexpectedness, the objective of describing an image in a witty manner often results in a trade-off between the description being witty and the description being relevant to the image. It may be interesting to study how the perceived wittiness of an image description varies as it includes more creative elements and becomes less relevant to the image. Another interesting factor that can influence perceived humor is presentation. For instance, the text in cartoons and memes are funny in their characteristic, informal font but may seem boring in other, more "serious" font.

Producing a description for an image that is perceived as witty is challenging because the description must achieve the fine balance between lending itself to easy resolution by the perceiver while not being impossible or too trivial. There are other challenges, however. For instance, automatic image recognition and captioning models, despite the great strides of advancement in recent times, are still imperfect. In our approach, these are cascading sources of error which could adversely affect the perceived wittiness of an image description.

In this work, we only consider the use of *words* that are perfect puns. Future work can extend our approach to explore the use of phrase-based and imperfect puns to create alternate interpretations of a sentence.

Our approach has no constraints on the modality of the input stimulus. It can be extended to generate witty responses to input stimuli of different modalities, e.g., text (to generate witty dialogue) or video (to generate witty video-descriptions).

6 Conclusion

We presented novel computational models to address the challenging task of producing contextually witty descriptions for a given image. We leverage linguistic wordplay, specifically puns in both retrieval and generation style models. We evaluate our models against meaningful baseline approaches via human studies. In a Turing test style evaluation, annotators find image descriptions from our generation model to be wittier than a human's witty description 55% of the time!

Acknowledgements

We thank Shubham Toshniwal for his help with the automatic speech recognition model. This work was supported in part by: NSF CAREER, ARO YIP, Google GFRA, ARL grant W911NF-15-2-0080, ONR grant N00014-16-1-2713, ONR YIP, Paul G. Allen Family Foundation award, and a Sloan Fellowship to DP; and NVIDIA GPU donations, Google GFRA, IBM Faculty Award, and Bloomberg Data Science Research Grant to MB.

References

Salvatore Attardo and Victor Raskin. 1991. Script theory revis (it) ed: Joke similarity and joke representation model. *Humor-International Journal of Humor Research* 4(3-4):293–348.

Kim Binsted. 1996. Machine humour: An implemented model of puns. *PhD Thesis, University of Edinburgh* .

Kim Binsted and Graeme Ritchie. 1997. Computational rules for generating punning riddles. *Humor: International Journal of Humor Research* .

Arjun Chandrasekaran, Ashwin Kalyan, Stanislaw Antol, Mohit Bansal, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2016. We are humor beings: Understanding and predicting visual humor. In *CVPR*.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, pages 248–255.

Preethi Jyothi and Karen Livescu. 2014. Revisiting word neighborhoods for speech recognition. *ACL 2014* page 1.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*. Springer, pages 740–755.

Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics - Volume 1*. Association for Computational Linguistics, ETMTNLP '02.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*.

Tristan Miller. 2016. *Adjusting Sense Representations for Word Sense Disambiguation and Automatic Pun Interpretation*. Ph.D. thesis, tuprints.

Tristan Miller and Iryna Gurevych. 2015. Automatic disambiguation of english puns. In *ACL (1)*. pages 719–729.

William J Pepicello and Thomas A Green. 1984. *Language of riddles: new perspectives*. The Ohio State University Press.

Sasa Petrovic and David Matthews. 2013. Unsupervised joke generation from big data. In *ACL*.

Dragomir Radev, Amanda Stent, Joel Tetreault, Aasish Pappu, Aikaterini Iliakopoulou, Agustin Chanfreau, Paloma de Juan, Jordi Vallmitjana, Alejandro Jaimes, Rahul Jha, et al. 2015. Humor in collective discourse: Unsupervised funniness detection in the new yorker cartoon caption contest. *arXiv preprint arXiv:1506.08126* .

Dafna Shahaf, Eric Horvitz, and Robert Mankoff. 2015. Inside jokes: Identifying humorous cartoon captions. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pages 1065–1074.

Oliviero Stock and Carlo Strapparava. 2005. HAHAcronym: A computational humor system. In *ACL*.

Jerry M Suls. 1972. A two-stage model for the appreciation of jokes and cartoons: An information-processing analysis. *The Psychology of Humor: Theoretical Perspectives and Empirical Issues* .

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 2818–2826.

Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2016. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence* .

William Yang Wang and Miaomiao Wen. 2015. I can has cheezburger? a nonparanormal approach to combining textual and visual information for predicting and generating popular meme descriptions. In *NAACL*.

Zhou Yu, Leah Nicolich-Henkin, Alan W Black, and Alex I Rudnicky. 2016. A wizard-of-oz study on a non-task-oriented dialog systems that reacts to user engagement. In *17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. page 55.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE International Conference on Computer Vision*. pages 19–27.

Arnold Zwicky and Elizabeth Zwicky. 1986. Imperfect puns, markedness, and phonological similarity: With fronds like these, who needs anemones. *Folia Linguistica* 20(3-4):493–503.